arXiv:1304.3868v1 [cs.DS] 14 Apr 2013

A 2-Approximation for Redundancy Aware Network Design with Laminar Demands

Siddharth Barman* Shuchi Chawla† Seeun Umboh‡

Abstract

In this note we study the redundancy aware network design problem with laminar demands. In this problem we are given a graph, several source-sink pairs, and a universe of data packets. Each sink desires a subset of the packets from its corresponding source. Our goal is to find a collection of paths, one for each source-sink pair, such that the total cost of routing packets over these paths is minimized. The cost of routing on an edge in the graph is proportional to the total size of the *distinct* packets that the edge carries. We assume that the collection of packet sets is laminar, that is, the packet sets desired by any two pairs of sources and sinks are either disjoint or one contains the other. For this setting, we present a primal-dual based 2-approximation, improving upon a logarithmic approximation due to Barman and Chawla [4].

*California Institute of Technology. barman@caltech.edu.
†University of Wisconsin – Madison. shuchi@cs.wisc.edu.
‡University of Wisconsin – Madison. seeun@cs.wisc.edu.

1 Introduction

In a network design problem we are given a graph and a traffic matrix specifying the amount of traffic to be sent between pairs of nodes in the graph. The goal is to build a minimum cost network that can support the given traffic matrix, as well as to specify a routing for the traffic on that network. The cost of routing traffic is a function of how much and what traffic each edge carries. For example, suppose that we can buy bandwidth on any edge at a fixed rate, then to minimize costs, we should route the traffic over shortest paths between the sources and the sinks. On the other hand, if the cost on an edge depends only on whether or not the edge carries traffic, but not on the amount of traffic it carries, then routing traffic along a minimum cost Steiner forest is the optimal solution.

We study a network design problem in the context of a communication network that can leverage the redundancy in data to save on routing costs. The input to the problem consists of multiple *commodities*, each with a source and several possible destinations that we collectively call terminals. Each commodity is composed of a number of different data packets drawn from a universe Π of packets; we call these sets of packets *demands*. Importantly, there is redundancy in traffic—different commodities may overlap in the sets of packets they contain, and so can benefit from using common routes. In particular, we assume that routers in the network can detect and remove duplicate packets, and can also replicate packets. This means that an edge carrying two different commodities that share common packets need only transmit each common packet once, thereby leading to savings in bandwidth and the cost of routing. Our goal, again, is to find a minimum cost routing for the given traffic matrix, assuming that we can buy bandwidth at a fixed rate on every edge. Formally, our solution specifies for each commodity a routing tree spanning all of the terminals for this commodity. The cost of this solution on any particular edge is proportional to the total size of the *distinct* data packets that the edge carries; in other words, it is a *coverage function* over the collection of commodities that use this edge. This problem was introduced in [4], and is called the *Redundancy Aware Network Design* problem (RAND).

Redundancy aware network design displays the same short-routes-versus-shared-routes tradeoff present in several classical network design problems with nonlinear costs, such as rent-or-buy network design [12, 9], access network design [2], and buy-at-bulk network design [3, 10, 13, 14]. However there are fundamental differences. The buy-at-bulk cost model is inspired by economies of scale in a physical commodity network—the volume of traffic that an edge carries is the sum of the volumes that the different commodities impose on it and the routing cost on the edge is a concave function of the total volume of traffic. On the other hand, in our setting, the volume of traffic itself is lowered due to the inherent nature of data traffic. In particular, this means that the savings achieved depend on the contents of the traffic and not just its quantity. We not only need to bundle traffic streams as much as we can, but we also need to decide the right sets of traffic streams to bundle. Consequently, the approximability of the problem also depends on the extent and manner in which different commodities share packets. When every source-sink pair in the network demands a distinct packet, that is, there is no data redundancy in the network, the problem reduces to finding the shortest route for each pair and can be solved exactly. When all of the demands are identical, the problem reduces to finding a single optimal Steiner forest over all of the terminal sets, and can be approximated within a factor of 2.

In this note we focus on a special case of the problem that captures much of its complexity—the *laminar demands* setting. In this setting the packet sets corresponding to the commodities form a laminar family: the packet sets of any two commodities are either completely disjoint or one contains the other. There is a natural heirarchy over commodities in this setting and any commodity can use for free an edge that is being used for another commodity that "dominates" it. So we may favor long routes for a commodity if they share

edges with a dominating commodity over short ones that do not share edges. Less intuitively, it may be useful to pick similar routes for two commodities with disjoint packets sets if a portion of the shared route can be used for a commodity that dominates both.

To form intuition for RAND in the laminar setting consider a setting with k different packets and $k+1$ commodities: for $i \leq k$ the demand set of commodity i contains only packet i, and demand set of commodity $k+1$ contains all of the k packets. Suppose also that every commodity has a single source and a single sink. Then, one approach to solving the problem is to first find a least cost path for commodity $k+1$, and then find least cost paths for the remaining commodities using the edges in the first path for free. This approach misses solutions where a slightly longer path for commodity $k+1$ is much more cost efficient for the remaining commodities than the shortest path for $k+1$. An alternative is to first find shortest paths for commodities 1 through k, and then find the least cost path for commodity $k+1$ that can use edges in previously picked paths at a cheaper cost. This misses solutions where picking slightly longer paths for commodities 1 through k leads to a greater sharing of the edges. The first approach is indeed the approach analyzed in [4] for the special case of the problem where there is a single source that belongs to all of the terminal sets. That paper shows that in any single source laminar setting routing commodities in order of decreasing sizes of demand sets achieves an $O(\log |\Pi|)$ approximation where $|\Pi|$ is the number of different packets in the universe, and this factor is tight.

In this paper we extend and improve the result of [4] to obtain a 2-approximation for the laminar demands setting with arbitrary terminal sets. Our approach is a hybrid of the two described above. At a high level, we first consider commodities in *increasing* order of the sizes of their demand sets. However, instead of committing to a single path for each commodity before considering the next, we keep around a collection of all possible near-optimal paths for the smaller demand sets before considering choices for the larger demand sets. Then in a second pass, we finalize a single path (tree) for each commodity, considering commodities in *decreasing* order of sizes of their demand sets. That is, we commit to paths for the larger demand sets before finalizing paths for the smaller demand sets. In order to maintain a collection of all near-optimal paths efficiently we use a primal-dual approach. The duals constructed for each commodity give a succinct description of all possible short paths connecting the source and the sink for that commodity. After having constructed all of the duals, we perform a reverse delete step that finalizes paths for commodities starting from the one with the largest demand and moving on to smaller demand sets.

1.1 Related work

The cost structure in RAND is uniform in the sense that costs on different edges are related through constant factors. Obtaining a randomized $O(\log n)$ approximation for network design problems with a uniform cost structure is often easy: we can use the tree embeddings of Bartal [5] and Fakcharoenphol et al. [6] to convert the graph into a distribution over trees such that distances between nodes are preserved to within logarithmic factors in expectation. Then the expected cost of the optimal routing over the (random) tree is related within logarithmic factors to the cost of the optimal routing over the graph. Moreover, the problem is easy to solve on trees, because there is a unique path between every pair of nodes.

As mentioned earlier, RAND is closely related but incomparable to other models of network design with uniform costs that display economies of scale. This includes, e.g., the uniform buy-at-bulk [3, 10, 13, 14], rent-or-buy [12, 9], and access network design [2, 8] problems. For all of these problems constant factor approximations are known in the uniform costs setting for the special case where all of the commodities share a common source. In the multi-commodity setting, i.e., with distinct sources and sinks, the rent-or-buy network design problem admits a 2-approximation [12, 9], but the buy-at-bulk network design problem is hard to approximate within poly-logarithmic factors [1].

In order to model information networks, Hayrapetyan et al. [11] considered a single-source network design problem in which the cost on an edge is a monotone submodular function of the commodities that use the edge. They obtain an $O(\log n)$ approximation via tree embeddings [5, 6], where n is the number of vertices in the graph. Note that the cost structure in RAND is a special case of the one in [11] (coverage functions are submodular). But, unlike [11], we consider arbitrary terminals sets and hence their setting does not generalize ours. In addition, we obtain a stronger approximation guarantee of 2.

1.2 Problem statement

The Redundancy Aware Network Design problem is defined as follows. We are given a graph $G = (V, E)$ with costs c_e on edges, a universe Π of packets, and g terminal sets $X_1, \ldots, X_g \subseteq V$. The demand set of terminal set X_j is denoted $D_j \subseteq \Pi$, and we denote the collection of all demand sets as $\mathcal{D}$. A solution consists of a collection of g Steiner trees $\mathcal{T} = \{T_1, \ldots, T_g\}$ where T_j is a Steiner tree spanning terminal set X_j. The trees specify how packets are to be routed over the edges: the packets of demand D_j are routed over edges of T_j. For a solution $\mathcal{T}$, the load on edge e is $\ell_e(\mathcal{T}) = |\bigcup_{i:e\in T_i} D_i|$, i.e. the total number of distinct packets being routed over edge e.[1] Our goal is to find a $\mathcal{T}$ so as to minimize the total cost $\sum_{e\in E} c_e \ell_e(\mathcal{T})$.

In the laminar demands setting, the collection of demand sets is laminar: for any $D, D' \in \mathcal{D}$, $D \cap D' \neq \emptyset$ implies that either $D \subseteq D'$ or $D' \subseteq D$.

2 A primal-dual 2-approximation

Our algorithm for the laminar demands case is an extension of the Goemans-Williamson primal-dual algorithm for the Steiner Forest Problem [7]. We begin by defining the primal and dual linear programs.

First, we transform our objective into a simpler form where the cost of each edge is charged to a collection of disjoint demand sets. In particular, given a solution $\mathcal{T}$, for an edge e consider the demand sets D that are maximal among the collection $\{D_j : e \in T_j\}$ of demand sets that this edge carries. Because of laminarity, these maximal demand sets are disjoint, and so the load on the edge is simply the sum of the sizes of these demand sets. Accordingly, let us define $H_D(\mathcal{T})$ to be the set of edges e such that D is a maximal set in $\{D_j : e \in T_j\}$. D will contribute to the load on these edges. Then, we can write the total cost of the solution $\mathcal{T}$ as

$$\ell(\mathcal{T}) = \sum_e c_e \ell_e(\mathcal{T}) = \sum_e \sum_{D:H_D(\mathcal{T})\ni e} c_e|D| = \sum_D |D| \sum_{e\in H_D(\mathcal{T})} c_e = \sum_D |D| c(H_D(\mathcal{T}))$$

Further note that in a feasible solution $\mathcal{T}$, for each commodity j, the subgraph $\bigcup_{D\supseteq D_j} H_D(\mathcal{T})$ spans the terminal set X_j because the edges of T_j carry a superset of D_j. Therefore, instead of specifying a Steiner tree for each terminal set, it suffices to specify a forest H_D for each demand set D such that each terminal set X_j is connected $\bigcup_{D\supseteq D_j} H_D$.

In the linear program below, the variable $x_{e,D}$ denotes whether $e \in H_D$. We denote by $\delta(S)$ the set of edges crossing a cut $S \subseteq V$, and $\mathcal{S}_D$ the collection of cuts $S \subseteq V$ that separates a terminal set X_j whose demand set D_j contains D. The cut constraints require that each terminal set X_j is connected by $\bigcup_{D\supseteq D_j} H_D$.

[1] More generally, we can define the load on an edge to be the total size of all of the distinct packets that an edge carries. Since our algorithm's performance or running time does not depend on the number of distinct packets, we may assume without loss of generality that all packets have equal size.

$$\begin{array}{ll} \text{minimize} & \displaystyle\sum_{e,D\in\mathcal{D}} x_{e,D}\cdot|D|c_e \\ \text{subject to} & \displaystyle\sum_{D'\supseteq D}\sum_{e\in\delta(S)} x_{e,D'}\geq 1 \qquad \forall D\in\mathcal{D}, S\in\mathcal{S}_D \end{array}$$

The corresponding dual linear program is as follows.

$$\begin{array}{ll} \text{maximize} & \displaystyle\sum_{D\in\mathcal{D},S\in\mathcal{S}_D} y_{D,S} \\ \text{subject to} & \displaystyle\sum_{D'\subseteq D}\sum_{S\in\mathcal{S}_{D'}:e\in\delta(S)} y_{D',S}\leq|D|c_e \qquad \forall e, D\in\mathcal{D} \end{array}$$

2.1 Algorithm

The algorithm starts with a dual ascent stage in which it adds edges to forests $\{F_D\}_{D\in\mathcal{D}}$, and ends with a pruning stage. In the following discussion, for a demand set $D\in\mathcal{D}$ we say that $S\in\mathcal{S}_D$ is a *D-unsatisfied cut* if $\bigcup_{D'\supseteq D} F_{D'}\cap\delta(S)=\emptyset$. We also say that an edge e is *D-tight* if

$$\sum_{D'\subseteq D}\sum_{S\in\mathcal{S}_{D'}:e\in\delta(S)} y_{D',S}=|D|c_e.$$

In the dual ascent stage, the algorithm raises duals in phases, one per demand set $D\in\mathcal{D}$ in order of increasing size. In phase D, while there exists a D-unsatisfied cut it alternates between raising duals of the minimal D-unsatisfied cuts and adding D-tight edges to F_D. We say that S is an *active set* in the current iteration of the inner while loop if it is a minimal D-unsatisfied cut. The algorithm ensures that at the end of phase D, the edges F_D are paid for by the dual and F_D is a Steiner forest for terminal sets whose demand set contains D. In the pruning stage, the algorithm processes the demand sets in order of decreasing size and removes unnecessary edges from $\{F_D\}_{D\in\mathcal{D}}$ and returns $\{H_D\}_{D\in\mathcal{D}}$.

We need the following lemma to prove that we can efficiently find active sets.

Lemma 1. *In any iteration in phase D, a set S is active if and only if it is a component of F_D and it separates a terminal set whose demand set contains D.*

Proof. Let S be an active set. By definition, S is a minimal cut in $\mathcal{S}_D$ such that $\bigcup_{D'\supseteq D} F_{D'}\cap\delta(S)=\emptyset$. Since $S\in\mathcal{S}_D$, it separates a terminal set whose demand set contains D. The algorithm processes the demand sets in increasing order of size, so we have $F_{D'}=\emptyset$ for $D'\supsetneq D$ and thus $F_D\cap\delta(S)=\emptyset$. This implies that $S\cap C=\emptyset$ or $S\cap C\supseteq C$ for every connected component C of F_D and so S is a superset of a union of connected components of F_D. By minimality, we have that S is a connected component of F_D.

For the converse, consider a connected component S' of F_D that separates a terminal set whose demand set contains D. By definition, we have $S'\in\mathcal{S}_D$. Since S' is a connected component of F_D and $F_{D'}=\emptyset$ for $D'\supsetneq D$, it is a minimal set in $\mathcal{S}_D$ such that $\bigcup_{D'\supseteq D} F_{D'}\cap\delta(S)=\emptyset$. Therefore S' is an active set. □

2.2 Analysis

Our analysis follows along the lines of the analysis for the Goemans-Williamson algorithm. We first prove that the primal and dual solutions generated by the algorithm are feasible.

Algorithm 1 Primal-Dual Algorithm for Laminar Buy-at-Bulk

Initialize $F_D \leftarrow \emptyset$ for all $D \in \mathcal{D}$ and $y_{D,S} \leftarrow 0$ for all $D \in \mathcal{D}, S \subseteq V$.
(Dual ascent stage)
for $D \in \mathcal{D}$ in increasing order of size **do**
  (Start of phase D)
  while there exists a D-unsatisfied cut **do**
    Simultaneously raise $y_{D,S}$ for active sets S until some edge e goes D-tight.
    $F_D \leftarrow F_D + e$.
  end while
  (End of phase D)
end for
(End of dual ascent stage)
(Pruning stage)
$H_D \leftarrow F_D$ for all $D \in \mathcal{D}$.
for $D \in \mathcal{D}$ in decreasing order of size **do**
  for $e \in H_D$ **do**
    if $(H_D - e) \cup \bigcup_{D' \supsetneq D} H_{D'}$ is a Steiner forest for terminal sets with demand set D **then**
      $H_D \leftarrow H_D - e$.
    end if
  end for
end for
(End of pruning stage)
return $\{H_D\}_D$

Lemma 2. *The primal solution $\{H_D\}_{D\in\mathcal{D}}$ and the dual solution $\{y_{D,S}\}_{D\in\mathcal{D},S\subseteq V}$ are feasible.*

Proof. We first prove that the primal solution is feasible. Consider an iteration during the pruning stage. We say that terminal set X_j is *H-disconnected* if it is disconnected with respect to edge set $\bigcup_{D\supseteq D_j} H_D$ and *H-connected* otherwise. We will show that all terminal sets are H-connected in all iterations of the pruning stage.

Observe that at the end of phase D, there are no D-unsatisfied cuts and $F_{D'} = \emptyset$ for $D' \supsetneq D$. Thus, all terminal sets with demand set D are connected with respect to edge set F_D. At the beginning of the pruning stage, we have $H_D = F_D$ for all $D \in \mathcal{D}$, and so all terminal sets are H-connected. Consider an iteration in which the algorithm deletes an edge e from H_D. By definition of H-disconnected, this can only cause a terminal set with demand set $D' \subseteq D$ to be H-disconnected. However, the algorithm will not delete e if it causes a terminal set with demand set D to be H-disconnected. Now consider a demand set $D' \subsetneq D$. Since $|D'| \leq |D|$, we still have $H_{D'} = F_{D'}$ so all terminal sets with demand set D' are H-connected. Thus, all terminal sets are H-connected throughout the pruning stage and so $\{H_D\}_{D\in\mathcal{D}}$ is a feasible primal solution.

The dual solution is feasible since the algorithm explicitly ensures that the dual variables in a tight constraint are not raised. □

Next, we prove that in each phase D of the dual raising stage, the current active sets has average degree with respect to edges $\bigcup_{D'\supseteq D} H_{D'}$ (formally defined below) at most 2 in every iteration. This in turn implies that the primal solution has cost at most twice the total dual value. Since the dual is feasible, we have that the algorithm gives a 2-approximation. We bound the average degree of active sets by showing that $\bigcup_{D'\supseteq D} H_{D'}$ is a forest and that no inactive set has degree 1.

Lemma 3. *For all $D \in \mathcal{D}$, we have that $\bigcup_{D'\supseteq D} H_{D'}$ is a forest.*

Proof. Suppose, towards a contradiction, that the statement is false. Let D be a maximal demand set such that $\bigcup_{D'\supseteq D} H_{D'}$ contains a cycle C. By maximality, there exists $e \in C \cap H_D$. Since e is in a cycle in $\bigcup_{D'\supseteq D} H_D$, we have that $(H_D - e) \cup \bigcup_{D'\supsetneq D} H_{D'}$ is still a Steiner forest for terminal sets with demand set D. Thus, the algorithm would have removed e from H_D and so we have a contradiction. □

For a subset of edges $E' \subseteq E$, let $\deg_{E'}(S) = |\delta(S) \cap E'|$ denote the number of edges in E' that leave S.

Lemma 4. *Consider an iteration in phase D of the dual raising stage. Let S be a connected component of F_D in this iteration. If $S \notin \mathcal{S}_D$, then $\sum_{D'\supseteq D} \deg_{H_{D'}}(S) \neq 1$.*

Proof. We prove the contrapositive. Suppose $\sum_{D'\supseteq D} \deg_{H_{D'}}(S) = 1$. Let e and $A \supseteq D$ be the unique edge and demand set, respectively, such that $e \in H_A \cap \delta(S)$. Since the algorithm did not delete e from H_A and $\bigcup_{D'\supseteq A} H_{D'}$ is acyclic by Lemma 3, there exists X_j with $D_j = A$ and $u, v \in X_j$ such that e is on the unique $u - v$ path in $\bigcup_{D'\supseteq A} H_{D'}$. Since $\sum_{D'\supseteq D} \deg_{H_{D'}}(S) = 1$, the path crosses S exactly once. Thus, we have that S separates u, v and so $S \in \mathcal{S}_A$. By definition of $\mathcal{S}_D$, we have $\mathcal{S}_A \subseteq \mathcal{S}_D$ and this completes the proof of the lemma. □

We are now ready to prove that the primal solution has cost at most twice the dual value.

Lemma 5. $\sum_D \sum_{e\in H_D} |D| c_e \leq 2 \sum_{D,S} y_{D,S}$.

Proof. Using the fact that we only add tight edges, we have

$$\begin{aligned}\sum_D \sum_{e \in H_D} |D| c_e &= \sum_D \sum_{e \in H_D} \left(\sum_{D' \subseteq D} \sum_{S \in \mathcal{S}_{D'} : e \in \delta(S)} y_{D',S} \right) \\ &= \sum_{D'} \sum_{S \in \mathcal{S}_{D'}} y_{D',S} \left(\sum_{D \supseteq D'} \sum_{e \in \delta(S) \cap H_D} 1 \right) \\ &= \sum_{D'} \sum_{S \in \mathcal{S}_{D'}} y_{D',S} \left(\sum_{D \supseteq D'} \deg_{H_D}(S) \right) \\ &= \sum_{D'} \sum_{S \in \mathcal{S}_{D'}} y_{D',S} \deg_{\bigcup_{D \supseteq D'} H_D}(S).\end{aligned}$$

The second equality is obtained by rearranging, and the last follows from the fact that each edge is in H_D for at most one $D \supseteq D'$.

Suppose that in an iteration in phase D', the dual for each active set is raised by Δ. Then we have that $\sum_{S \in \mathcal{S}_{D'}} y_{D',S} \deg_{\bigcup_{D \supseteq D'} H_D}(S)$ increases by $\Delta \cdot \sum_{S \text{ active}} \deg_{\bigcup_{D \supseteq D'} H_D}(S)$, and $\sum_{D,S} y_{D,S}$ increases by $\Delta \cdot$ # of active sets. So it suffices to prove that in each phase D' and in each iteration within the phase, the average degree of active sets is at most 2:

$$\sum_{S \text{ active}} \deg_{\bigcup_{D \supseteq D'} H_D}(S) \leq 2 \cdot \text{\# of active sets.}$$

Fix an iteration in phase D'. Note that each active set correspond to some connected component of $F_{D'}$ by Lemma 1. Let G' be a graph whose nodes are connected components of $F_{D'}$ and whose edge set is $\bigcup_{D \supseteq D'} H_D$. The degree of a node in G' is equal to the degree of the corresponding set with respect to edge set $\bigcup_{D \supseteq D'} H_D$. Let us say that a node of G' corresponding to an active set is an *active node*, and that any other node is *inactive*. We want to show that the average degree of active nodes in G' is at most 2. Suppose we remove all isolated nodes from G'. In the resulting graph, the degree of each inactive node is at least 2 by Lemma 4, and the average degree is at most 2 by Lemma 3. Thus, the average degree of active nodes is at most 2. □

Lemmas 2 and 5 gives us the following theorem.

Theorem 6. *Algorithm 1 is a 2-approximation for redundancy aware network design with laminar demands.*

References

[1] Matthew Andrews. Hardness of buy-at-bulk network design. In *Foundations of Computer Science, 2004. Proceedings. 45th Annual IEEE Symposium on*, pages 115–124. IEEE, 2004.

[2] Matthew Andrews and Lisa Zhang. Approximation algorithms for access network design. *Algorithmica*, 34(2):197–215, 2002.

[3] Baruch Awerbuch and Yossi Azar. Buy-at-bulk network design. In *Proceedings of the 38th IEEE Symposium on Foundations of Computer Science*, pages 542–547, 1997.

[4] Siddharth Barman and Shuchi Chawla. Traffic-redundancy aware network design. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '12, pages 1487–1498. SIAM, 2012.

[5] Y. Bartal. Probabilistic approximation of metric spaces and its algorithmic applications. In *Proceedings of the 37th Annual Symposium on Foundations of Computer Science*, FOCS '96, pages 184–, Washington, DC, USA, 1996. IEEE Computer Society.

[6] Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 448–455, 2003.

[7] M. Goemans and D. Williamson. A general approximation technique for constrained forest problems. *SIAM Journal on Computing*, 24(2):296–317, 1995.

[8] S. Guha, A. Meyerson, and K. Munagala. Hierarchical placement and network design problems. In *Foundations of Computer Science, 2000. Proceedings. 41st Annual Symposium on*, pages 603–612, 2000.

[9] Anupam Gupta, Amit Kumar, Martin Pál, and Tim Roughgarden. Approximation via cost-sharing: A simple approximation algorithm for the multicommodity rent-or-buy problem. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '03, pages 606–617, 2003.

[10] Anupam Gupta, Amit Kumar, and Tim Roughgarden. Simpler and better approximation algorithms for network design. In *STOC '03: Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 365–372, 2003.

[11] Ara Hayrapetyan, Chaitanya Swamy, and Éva Tardos. Network design for information networks. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 933–942. Society for Industrial and Applied Mathematics, 2005.

[12] Amit Kumar, Anupam Gupta, and Tim Roughgarden. A constant-factor approximation algorithm for the multicommodity rent-or-buy problem. In *Proceedings of the 43rd Symposium on Foundations of Computer Science*, FOCS '02, pages 333–344, 2002.

[13] Adam Meyerson, Kamesh Munagala, and Serge Plotkin. Cost-distance: Two metric network design. *SIAM J. Comput.*, 38(4):1648–1659, December 2008.

[14] Kunal Talwar. The Single-Sink Buy-at-Bulk LP Has Constant Integrality Gap. In *Proceedings of the 9th International IPCO Conference on Integer Programming and Combinatorial Optimization*, pages 475–486, 2002.